

The CloudVeneto initiative: 10 years of operations to support interdisciplinary open science

Paolo Andreetto¹, Fulvia Costa¹, Alberto Crescente¹, Sergio Fantinel², Federica Fanzago^{1,1}, Alberto Garfagnini⁴, Loris Lazzaro⁴, Rita Lenzo¹, Paolo Emilio Mazzon³, Matteo Menguzzato⁴, Gianpietro Sella⁵, Massimo Sgaravatto¹, Sergio Traldi¹, Marco Verlato¹ and Lisa Zangrando¹

¹INFN, Sezione di Padova, Via Marzolo 8, 35131 Padova, Italy

²INFN, Laboratori Nazionali di Legnaro, Viale dell'Università 2, 35020 Legnaro (Padova), Italy

³Padova Neuroscience Center, Università di Padova, Via Orus 2/B, 35131 Padova, Italy

⁴Dipartimento di Fisica e Astronomia "Galileo Galilei", Università di Padova, Via Marzolo 8, 35131 Padova, Italy

⁵Dipartimento di Scienze Chimiche, Università di Padova, Via Marzolo 1, 35131 Padova, Italy

Abstract. CloudVeneto is a private cloud targeted to scientific communities based on OpenStack software. It was designed in 2013 and put in operation one year later, to support INFN projects, mainly HEP ones. Its resources are physically distributed among two sites: the Physics Department of University of Padova-INFN Padova Unit and the INFN Legnaro National Laboratories. During these 10 years CloudVeneto evolved to integrate also resources funded by ten Departments of the University of Padova, and to support several scientific disciplines of different domains. The use cases the communities have to face up often show a common pattern. This was an opportunity for us to develop and improve the services on our infrastructure to provide common solutions to different use cases. It happened for example with the Container as a Service (CaaS) that makes the management of Kubernetes clusters easier from a user point of view. Moreover, CloudVeneto joined the INFN national cloud infrastructure (INFN Cloud), making available some resources to this federated infrastructure. CloudVeneto is also involved in an R&D project to realize a distributed analysis facility for the CMS experiment based on the HTCondor batch system. In this paper we describe some use-cases of different projects pointing out the common patterns and the new implementations and configurations done in the infrastructure.

1 Introduction

The CloudVeneto infrastructure has been widely described in two previous articles [1,2]. In section 2 we'll quickly summarize its current status and capacity, and its recent integration with the INFN Cloud federation. In section 3 we'll describe a newly developed service (CloudVeneto-CaaS - Container as a Service) providing a fully managed orchestration platform of containers based on Kubernetes as a "central" cloud service. Its use by two scientific

¹ Corresponding author: Federica.Fanzago@pd.infn.it

communities served by CloudVeneto is described in sections 4 and 5. In section 6 we'll illustrate an example in the Biomedical Sciences field showcasing another common pattern usage based on clusters of VMs (hundreds of VCPUs) managed by batch systems (typically HTCondor, Sun Grid Engine and Slurm) to process data. In section 7 we'll focus on the involvement of CloudVeneto in the INFN CMS Analysis Facility R&D project. Section 8 concludes the article.

2 An update of the CloudVeneto infrastructure

CloudVeneto infrastructure is currently supporting about 500 users and 100 scientific research projects lead by INFN and University of Padova user communities. Fig. 1 shows the user growth rate since April 2020 (+80%).



Fig. 1. Number of users registered to CloudVeneto in the last three years.

It is based on the OpenStack software (Yoga version is currently installed) and many other open source tools and technologies allowing to implement basic compute and storage services, higher level services, deployment automation, monitoring, accounting, security incident handling, etc.

INFN Padova data center and INFN Legnaro National Laboratories, geographically separated by 10 km, host the nodes and the needed hardware equipment connected with a dedicated 10 Gbps optical network link. All OpenStack services are deployed ensuring high availability and fault tolerance making use of HAProxy/Keepalived, RabbitMQ and Percona XtraDB clusters. A unique system, implemented through Ceph [3], is used to provide storage for different clients and for different use cases. It provides block storage to: image (glance) and volume (cinder) OpenStack services; a Proxmox cluster hosting several ancillary services. It also provides, through a highly available configuration, object storage, exposed through S3 and swift. Different hardware classes and customized pool configurations (replicated pools, erasure code pools) can address multiple use cases. A dedicated network setup was configured to allow cloud instances to efficiently access the data of LHC experiments stored in the dCache instance of the LNL-Padova Tier2.

The hardware resources have been renewed and grown throughout the years with the funds provided by INFN and ten University of Padova Departments. Fig. 2 shows the growth over time of the compute cores since the beginning of the operations.

Table 1 shows in detail the total capacity of storage and computing resources achieved in 2023, funded by INFN and University of Padova administrations.




Fig. 2. Infrastructure growth in terms of number of cores (in HyperThreading) in the last ten years.

Table 1. CloudVeneto storage and computing resources by owner.

Owner	Storage (TB, raw)	# Compute Nodes	# Cores (in HT)	RAM (GB)	# GPUs
INFN	943	41	2144	10336	13
University	440	24	2072	10400	12
Total	1383	65	4216	20736	25

Some KVM hypervisors are configured with PCI passthrough virtualization to support 25 GPGPU cards of different models: 2 NVIDIA TITAN Xp, 1 NVIDIA Quadro RTX 6000, 1 NVIDIA GeForce GTX TITAN, 12 NVIDIA Tesla T4, 4 NVIDIA Tesla V100, 2 NVIDIA A30, 2 NVIDIA A2, 1 NVIDIA RTX A4000.

In the last three years the CloudVeneto team has been involved in the design and implementation of the INFN Cloud [4]. It is a federation of distributed infrastructures with the target to exploit all the computational resources available in the INFN and to implement a “national cloud infrastructure” for the research [5]. Its services are based on modular

components [6], covering IaaS, PaaS and SaaS and can be deployed in a simple way in the federated resources. CloudVeneto is fully compliant with the INFN Cloud “rules of participation”, and part of CloudVeneto resources are federated in INFN Cloud. Authentication of INFN Cloud users is implemented through the INDIGO Identity and Access Management (IAM) service [7]. Federated users are then mapped locally to dedicated OpenStack projects.

3 CloudVeneto-CaaS - Container as a Service

The Container as a Service (CaaS) represents a model for delivering containerized software packages via cloud resources. Unlike the more well-known Kubernetes as a Service (KaaS), where users are responsible for creating and managing their own Kubernetes cluster, CaaS offers a fully managed orchestration platform as a cloud service.

CaaS offers multiple advantages compared to KaaS, benefiting both users and service providers. Users don't need to have in-depth extensive knowledge of Kubernetes infrastructure management; they can simply hand over their container to the cloud provider using a web interface or command line interface. On the provider side, CaaS avoids the proliferation of individual Kubernetes clusters for each user, which can lead to unnecessary consumption of cloud resources when these clusters remain inactive for extended periods. This is very appreciated for IaaS with a limited amount of resources as CloudVeneto. The only drawback is that CaaS users don't have their private Kubernetes cluster and are limited in resource utilization, sharing resources with other users, and administrative capabilities.

Our CloudVeneto-CaaS solution aims to combine the benefits of both paradigms, offering the ease of use of CaaS while also providing users with the functionality to own private cloud resources for running containers. Users can create their Kubernetes nodes and integrate them into a centralized Control Plane, shared with other users but administered by CloudVeneto operators. Users have an integrated interface to request the on-demand deployment of their node sets, effectively creating logical clusters. They have control over cluster size and properties (CPU, RAM, storage) for individual nodes but delegate technical aspects like deployment, configuration, updates, and monitoring to CaaS. This model relieves users from the administrator role while retaining the advantages of a private cluster, even within a multi-user context.

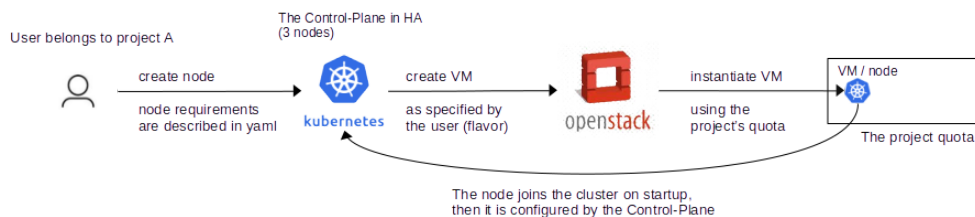


Fig. 3. Simplified schema of the CloudVeneto-CaaS architecture showing the steps allowing the users to create a Kubernetes node on the OpenStack project resources where they are authorized.

The sharing of our CaaS among CloudVeneto users is made possible through native Kubernetes support for multi-tenancy. This architectural approach groups users into CloudVeneto projects and uses namespaces to isolate groups of resources and users within the same cluster through policies and access rules (Role-Based Access Control, RBAC). While isolation isn't guaranteed between users of the same namespace, this organization facilitates resource allocation based on CloudVeneto projects, preventing resource waste. This allocation

strategy in fact prevents the prolonged idle state of cloud resources. Essentially, we consider that it's more efficient and equitable to restrict users to consume only the resources allocated to their respective CloudVeneto projects, rather than providing them with a shared pool of resources that could lead to contention and inefficiency.

To overcome specific limitations related to Kubernetes isolation and security capabilities, we've introduced integrated add-ons. These enhancements encompass authentication mechanisms based on Keystone and IAM tokens, along with refined authorization procedures to guarantee comprehensive user and resource isolation. Furthermore, we've integrated an operator into our system, specifically designed for the purpose of creating, configuring, and overseeing Kubernetes nodes on cloud infrastructures. This operator's functionality is particularly tailored to the OpenStack-based setup utilized by CloudVeneto, as shown in Fig. 3.

4 Applied research in Nuclear Physics use cases

In scientific research, efficiency and flexibility in handling complex computational tasks are crucial. CloudVeneto-CaaS significantly improved the accessibility of CloudVeneto resources for projects that work in the field of nuclear physics, like SPES [8] and ISOLPHARM [9].

The SPES project is a facility for the production of radioactive ion beams (RIBs) that will operate in the next few years at the INFN Legnaro National Laboratories. Recently, extensive calculations were performed in CloudVeneto with FLUKA Monte Carlo code, to evaluate the residual radioactivity generated on the materials of the SPES production apparatus in different times after the beginning of the facility operation, and specifically to estimate the external exposure of the personnel operating in the facility [10,11]. These studies must necessarily be addressed in the design phase of the plant and must include all the operation phases, up to the final disposal of the facility. Therefore, a powerful cloud-based computing resource was needed to reach a significant statistical precision, especially when simulating the dose/months or dose/years after the start of apparatus operations, considering different radiation scenarios.

Similarly, GEANT4 was used to perform simulations in CloudVeneto for the purposes of the ISOLPHARM project, an application of the SPES project devoted to the development of novel radiopharmaceuticals based on ^{111}Ag radioisotope core. A wide campaign of simulations was realized to evaluate the performance of a future ^{111}Ag production within the SPES facility [12]. More recently a study of the production and characterization of ^{111}Ag in a TRIGA Mark II nuclear research reactor was made possible thanks to the availability of CloudVeneto infrastructure [13].

To fulfil these goals, CloudVeneto-CaaS streamlined tasks such as provisioning VMs and deploying Kubernetes clusters, allowing researchers to efficiently perform Monte Carlo simulations within the SPES facility and for the ISOLPHARM project. CloudVeneto-CaaS enabled parallel execution of complex simulations, generating a substantial number of events in a short time frame, thanks to the optimized logical cluster managed by the project. In fact, in the last three years CloudVeneto-CaaS has simulated approximately a total of 100 billion events through 10 thousand parallel jobs.

5 Quantum Computing research use case

In the realm of advanced quantum research, the Quantum Tensor Network Emulation and Applications (QuantumTEA) project [14,15] enables the emulation of quantum circuits, streamlining the development and evaluation of quantum machine learning algorithms at a significant scale, capable of accommodating a high number of qubits [16]. Subsequently, these algorithms can be executed on dedicated quantum processors.

Notably, the project is planning to harness CloudVeneto computing resources and services, specifically the CloudVeneto-CaaS platform, to run quantum circuits on dedicated computing nodes through a newly developed Kubernetes operator designed explicitly for this purpose. This operator, named QuantumTEA-operator, is currently under test and aims at enhancing Kubernetes capabilities by incorporating specific functionalities required by the QuantumTEA project.

6 Biomedical Sciences use case

Another solution common to different scientific communities is the use of clusters of VMs managed by batch systems like HTCondor, Sun Grid Engine (SGE), Slurm, that share through NFS a filesystem backed on cloud storage. It is for example the case of a research group of the Department of Biomedical Sciences. They used CloudVeneto resources to update the entries of MobiDB [17] and RepeatsDB [18] data banks.

The MobiDB database is a knowledge base of intrinsically disordered proteins. In order to update it, the MobiDB-lite disorder predictor software [19] runs in parallel on 24 VMs with 8 VCPUs each, managed by the SGE batch system, in order to process 220 million protein sequences from the UniProtKB/TrEMBL data bank. Each iteration processes 1.5 million sequences reading 500 MB of input data and producing about 6 GB of output data in 11 hours over 192 VCPUs. Both input and output data are stored in cloud Ceph volumes of a few TBs mounted on the master VM and exported via NFS to the slave VMs.

The RepeatsDB database provides annotations and classification for protein tandem repeat structures from the Protein Data Bank (PDB). In this case the input data of a typical updating job amounts to about 100 GB and their parallel processing generates 5 GB of output data in 21 hours, always over 192 VCPUs.

In both cases, the generated output is copied from the master VM to the local servers of the research group that are connected through a VPN. From there are then imported in the MobiDB and RepeatsDB data banks.

7 CloudVeneto and the INFN CMS Analysis Facility

CloudVeneto is also involved in the INFN CMS Analysis Facility R&D project (AF) [20]. The INFN Analysis facility is conceived as a testbed for analysis that can be used in batch mode or for interactive analyses also using new frameworks.

For decades analysis workflows of LHC experiments have been tested locally using small subsets of data and only later, if tests were successfully executed, submitted in batch mode to remote clusters accessing complete datasets. The batch model implemented over the grid has shown some limitations related to the efficiency in resource usage and to the time needed to get results. Furthermore, the planned upgrade to HL-LHC will produce an amount of data orders of magnitude higher than we have to manage today, and more complex. The grid model will not be able to provide enough computing power for their management and analysis. For these reasons the exploitation of new kinds of resources (e.g provided by clouds, HPCs and FPGAs) and of new analysis models (e.g. using DataFrame) are steps to be evaluated and implemented.

This R&D project wants to be a platform to allow these kinds of tests. Thanks to the evolution of technologies and softwares, the interactive analysis can be run in a simpler way and the switch to the batch mode can be more transparent. Dask, Jupyter, Kubernetes and HTCondor are among the enabling technologies used in this Analysis Facility that allow to

abstract the physics analysis workflow from resources. The usage of new technologies is strongly recommended for Run3 and beyond.

The AF wants to demonstrate the possibilities to run interactive analysis applying big data analytics methodology and to realize a national pool, independent from heterogeneous resources, open to run analysis in a batch legacy way. Initial activities are using both a small fraction of computing resources at some Tier2 and some nodes made available in CloudVeneto. Central services of AF (e.g. JupyterHub and Central manager service of HTCondor) are installed on the INFN Cloud backbone, in a Kubernetes cluster.

The user entry point is provided by an JupyterHub portal that allows users to start their own JupyterLab interface. User authentication and authorization are provided by the INDIGO IAM service through CERN Single Sign-On (SSO) after accepting the Acceptable Usage Policy (AUP) terms. To instantiate the JupyterLab server, images are available already configured to provide an environment where to perform interactive analysis or just submit jobs to a pool of resources in a cluster. Users can also access the batch system via command line. The batch system is implemented via HTCondor that allows to integrate heterogeneous resources geographically distributed.

Dask is the tool that makes it easy to parallelize any Python code and to scale the Python libraries to a cluster. To improve the interactivity, HTCondor allows the distribution of Dask nodes over its pool. The software of the experiment is available on remote resources via CVMFS and data is accessed via XRootD service.

Currently CloudVeneto provides 12 VMs for this testbed with a total of 96 VCPUs and 3 TB of disk storage, and a XrootD server with 32 VCPUs and 10 TB of disk storage.

Each virtual machine is configured, via docker containers, as an HTCondor node. These nodes are part of a HTCondor pool that currently includes resources provided by some Italian CMS Tiers.

8 Conclusions

This article describes the current status of the 10 years old CloudVeneto infrastructure, one of the first OpenStack based IaaS cloud offerings put in production at INFN and then evolved and supported by a joint effort of INFN and the University of Padova.

During its lifetime new higher level services have been developed to face the challenges posed by the continuously growing number of user communities adopting the cloud computing model for their daily research activity. In particular, the CloudVeneto-CaaS service helped the Applied Nuclear Physics and the Quantum Computing communities to adopt a cloud native computing model without the burden to administrate complex container orchestration systems, while clusters of dozens of VMs managed by traditional batch systems instantiated on demand and enabled with new container based technologies and softwares allow large scale data processing as well as interactive analysis for Biomedical and HEP communities without the limitations and the poor flexibility of the grid model.

Eventually, CloudVeneto has been successfully federated with INFN Cloud, the INFN national distributed computing infrastructure. This provides the INFN users with a further set of modular services at PaaS and SaaS level that can be deployed on the federated resources in an easy and automated way.

References

1. P. Andreetto et al., *Merging OpenStack based private clouds: the case of CloudVeneto.it*, EPJ Web of Conferences **214**, 07010 (2019), doi: 10.1051/epjconf/201921407010
2. P. Andreetto et al., *Evolution of the CloudVeneto.it private cloud to support research and innovation*, EPJ Web of Conferences **245**, 07013 (2020), doi: 10.1051/epjconf/202024507013
3. Ceph home page, <https://ceph.io/>
4. INFN Cloud - Cloud resources for research, <https://www.cloud.infn.it>
5. M. Antonacci et al., *INFN and the evolution of distributed scientific computing in Italy*, Proceedings of CHEP2023
6. M. Antonacci et al., *A dynamic and extensible web portal enabling the deployment of scientific virtual computational environments on hybrid e-infrastructures*, 14th International Workshop on Science Gateways (IWSG 2022), 15-17 June 2022, doi: 10.5281/zenodo.7883082
7. A. Ceccanti et al., *Beyond X.509: token-based authentication and authorization for HEP*, EPJ Web of Conferences **214**, 09002 (2019), doi: 10.1051/epjconf/201921409002
8. G. Prete et al., *The SPES project at the INFN- Laboratori Nazionali di Legnaro*, EPJ Web of Conferences **66**, 11030 (2014), doi: 10.1051/epjconf/20146611030
9. S. Corradetti et al., *The ISOLPHARM Project: Production of High Specific Activity Medical Radioisotopes at SPES*, Nuclear Physics News **30** (2020) 33, doi: 10.1080/10619127.2020.1752100
10. A. Donzella et al., *Residual activation of the SPES Front-End system: a comparative study between the MCNPX and FLUKA codes*, Eur. Phys. J. A (2020) 56:54, doi: 10.1140/epja/s10050-020-00068-1
11. A. Monetti, *Study of the radionuclide deposition in the radioactive ion line of the Selective Production of Exotic Species (SPES) facility*, submitted to Applied Radiation and Isotopes
12. M. Ballan et al., *Preliminary evaluation of the production of non-carrier added ¹¹¹Ag as core of a therapeutic radiopharmaceutical in the framework of ISOLPHARM_Ag experiment*, Applied Radiation and Isotopes **164** (2020) 109258, doi: 10.1016/j.apradiso.2020.109258
13. L. Morselli et al., *Production and characterization of ¹¹¹Ag radioisotope for medical use in a TRIGA Mark II nuclear research reactor*, Applied Radiation and Isotopes **197** (2023) 110798, doi: 10.1016/j.apradiso.2023.110798
14. Quantum Tea home page, <https://www.quantumtea.it/>
15. D. Jaschke et al., *Ab-initio two-dimensional digital twin for quantum computer benchmarking*, arXiv preprint arXiv:2210.03763 (2022), doi: 10.48550/arXiv.2210.03763
16. M. Ballarin et al., *Entanglement entropy production in Quantum Neural Networks*, Quantum **7**, 1023 (2023), doi: 10.22331/q-2023-05-31-1023
17. D. Piovesan et al., *MobiDB: 10 years of intrinsically disordered proteins*, Nucleic Acids Research, Volume 51, Issue D1, 6 January 2023, Pages D438–D444, doi: 10.1093/nar/gkac1065.
18. L. Paladin et al., *RepeatsDB in 2021: improved data and extended classification for protein tandem repeat structures*, Nucleic Acids Research, Volume 49, Issue D1, 8 January 2021, Pages D452–D457, doi: 10.1093/nar/gkaa1097
19. M. Necci et al., *MobiDB-lite 3.0: fast consensus annotation of intrinsic disorder flavors in proteins*, Bioinformatics, Volume 36, Issue 22-23, December 2020, Pages 5533–5534, doi: 10.1093/bioinformatics/btaa1045
20. INFN CMS Analysis Facility, <https://infn-cms-analysisfacility.readthedocs.io/>